Irreversible Local Markov Chains with Rapid Convergence towards Equilibrium

Sebastian C. Kapfer[1,*] and Werner Krauth[2,3,†]

[1]Theoretische Physik 1, FAU Erlangen-Nürnberg, Staudtstraße 7, 91058 Erlangen, Germany

[2]Laboratoire de Physique Statistique, Département de physique de l'ENS, Ecole Normale Supérieure, PSL Research University, Université Paris Diderot, Sorbonne Paris Cité, Sorbonne Universités, UPMC Université Paris 06, CNRS, 75005 Paris, France

[3]Department of Physics, Graduate School of Science, The University of Tokyo, 7-3-1 Hongo, Bunkyo, Tokyo, Japan

(Received 25 May 2017; revised manuscript received 5 October 2017; published 15 December 2017)

We study the continuous one-dimensional hard-sphere model and present irreversible local Markov chains that mix on faster time scales than the reversible heat bath or Metropolis algorithms. The mixing time scales appear to fall into two distinct universality classes, both faster than for reversible local Markov chains. The event-chain algorithm, the infinitesimal limit of one of these Markov chains, belongs to the class presenting the fastest decay. For the lattice-gas limit of the hard-sphere model, reversible local Markov chains correspond to the symmetric simple exclusion process (SEP) with periodic boundary conditions. The two universality classes for irreversible Markov chains are realized by the totally asymmetric SEP (TASEP), and by a faster variant (lifted TASEP) that we propose here. We discuss how our irreversible hard-sphere Markov chains generalize to arbitrary repulsive pair interactions and carry over to higher dimensions through the concept of lifted Markov chains and the recently introduced factorized Metropolis acceptance rule.

DOI: 10.1103/PhysRevLett.119.240603

The hard-sphere model plays a central role in statistical mechanics. In three spatial dimensions (3D), the classical hard-sphere crystal melts in a first-order phase transition [1], whereas 2D hard spheres undergo a sequence of two phase transitions that have been characterized only recently [2,3]. Hard spheres have established paradigms for order-from-disorder phenomena driven by the depletion interaction [4,5] and for 2D melting with its dissociation of orientational and positional order [6]. The dynamics of the hard-sphere model has also been the focus of great attention, from the first algorithmic implementation of Newtonian mechanics through event-driven molecular dynamics [7] and the discovery of algebraically decaying velocity autocorrelations [8] to insights into the glass transition [9] and granular materials [10], and from the first definition of Markov-chain Monte Carlo dynamics [11] to rigorous convergence rates towards equilibrium in some special cases [12].

In 1D, the thermodynamics and the static correlation functions of finite hard-sphere systems can be computed exactly [5,13]. Newtonian dynamics is pathological for equal sphere masses, because colliding spheres simply exchange their velocities without mixing them. Stochastic dynamics, however, may converge to equilibrium. For example, reversible heat bath dynamics mixes (that is, converges towards equilibrium from an arbitrary starting configuration) in at most $\sim N^3 \log N$ individual steps for N spheres [14].

In the present Letter, we study irreversible local Markov chains for 1D hard-sphere systems that violate the detailed-balance condition yet still converge towards equilibrium. We show, by numerical simulation, that these irreversible Markov chains typically mix faster than reversible ones, and that they fall into two universality classes (see the Supplemental Material Section I [15] for background on balance conditions and the Supplemental Material Section II [15] for details on mixing and correlation times). The first one mixes in $O(N^{5/2})$ steps, and is related to the totally asymmetric simple exclusion process (TASEP, see [16–18]). The other mixes in $O(N^2 \log N)$ and comprises the event-chain algorithm (ECMC) [19] and a modified lifted TASEP that we propose in this Letter. We refer to it as the lifted TASEP class. The framework for our approach to irreversible Markov chains is provided by the lifting concept [20,21] together with a factorized Metropolis acceptance rule [22] which considerably extend the range of applications for irreversible Markov chains (see Ref. [23] for a review). Some evidence for reduced mixing time scales has already been obtained [24].

In order for a reversible or irreversible Markov chain to converge to the thermodynamic equilibrium given by π, the total probability flow $\mathcal{F}_a$ into a configuration a must satisfy the global balance condition

$$\mathcal{F}_a \equiv \sum_b \pi(b) T(b \to a) = \pi(a), \tag{1}$$

where $T(b \to a)$ is the algorithmic transition probability from b to a. In the following, we distinguish between "accepted" flow $\mathcal{A}(b \to a) = \pi(b) T(b \to a)$ from configurations $b \neq a$ to a and "rejected" flow $\mathcal{R} = \pi(a) T(a \to a)$ which results from an attempted move from a that was not accepted. The global balance condition enforces stationarity of π under multiplication with the transfer matrix T (see the Supplemental Material Section I [15] for definitions). The special condition realized in reversible Markov chains

is the detailed balance $\pi(b)T(b \to a) = \pi(a)T(a \to b)$ which, in terms of the probability flows, is simply $\mathcal{A}(a \to b) = \mathcal{A}(b \to a)$, and which implies Eq. (1).

For concreteness, we restrict ourselves to N hard spheres of diameter d on a circle of length L, so that the free space is $L_{\text{free}} = L - Nd$, and the mean gap between spheres is $\ell_{\text{free}} = L_{\text{free}}/N$. All valid configurations a have the same statistical weight $\pi(a) = 1$. They consist in ordered particle positions $a = (\ldots, x_{i-1}, x_i, x_{i+1}, \ldots)$ with gap variables $\delta_i = x_i - x_{i-1} - d \geq 0$ and appropriate periodic boundary conditions. The partition function $Z \sim (L_{\text{free}})^N$ is analytic for all densities in the thermodynamic limit, and no phase transition takes place [5,13]. The model is isomorphic to N point particles on a circle of length L_{free} with the same gap variables and an interaction $V(\delta_i < 0) = \infty$ and $V(\delta_i \geq 0) = 0$ implementing both the nonoverlap and the ordering constraint. Because of this mapping onto a gas of free particles with mean gap ℓ_{free}, and because the step size distribution of our Markov chains scales with ℓ_{free}, their dynamics and, in particular, their mixing times do not depend on density (see the Supplemental Material Section III [15] for details). Nevertheless, the spatial correlation length of the hard-sphere model diverges in the close-packing limit.

First, we consider the reversible heat bath algorithm, which moves at each time step $t = 0, 1, \ldots$ a random sphere i to a random position between spheres $i-1$ and $i+1$ [x_i is uniformly sampled in $(x_{i-1} + d, x_{i+1} - d)$]. When studying the mixing dynamics, we ignore trivial uniform rotations of the configuration (which only mix in $\sim N^4$ steps [14]). Thus, we restrict our attention to quantities that can be expressed in terms of the δ_i and focus on the slow, large-scale density fluctuations. The reversible heat bath algorithm is known to mix in at most $\sim N^3 \log N$ steps. We assume that the slowest time scale is exposed by tracking the distribution $\pi(u_i)$ of any half-system distance $u_i = \delta_i + \delta_{i+1} + \cdots + \delta_{i+N/2}$ from a compact initial configuration at $t = 0$, where the variance of u_i equals $\text{Var}u_i = {L_{\text{free}}}^2/4$, towards equilibrium at $t \sim \tau_{\text{mix}}$, where $\text{Var}u_i = {L_{\text{free}}}^2/(4N + 4)$ (see the Supplemental Material Section II [15] for details). Our simulations, indeed, show, in agreement with the rigorous bounds [14], that $\sim N^3$ steps are insufficient for mixing [see Fig. 1(a)], while $\sim N^3 \log N$ steps suffice [see Fig. 1(b)]. The diverging slope of $\text{Var}u_i$ at τ_{mix} signals the cutoff phenomenon [25]. Our method also recovers the correct mixing time for the related discrete symmetric simple exclusion process (SEP) model (see below). For this model, the leading term of τ_{mix} is known rigorously [26] and scales as $N^3 \log N$ (notwithstanding the absence of the logarithm in the spectral gap). Thus, our numerical method reliably detects mixing times, including logarithmic corrections and prefactors.

Plots analogous to Figs. 1(a) and 1(b) obtain the mixing time scales for all Markov chains studied in the present Letter [see Fig. 1(c)]. For the heat bath algorithm, a simple

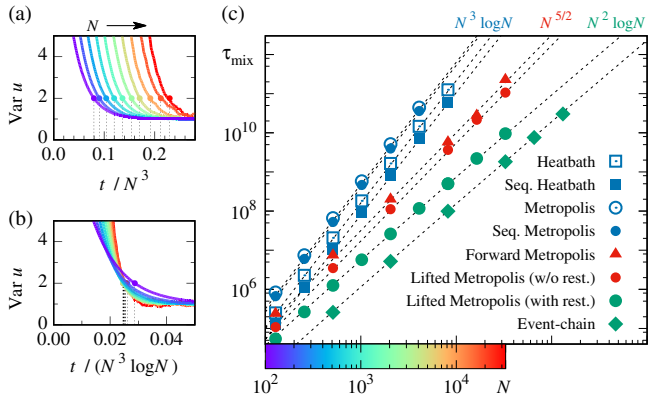

FIG. 1. Mixing of local 1D Markov chains. (a) Relaxation of $\text{Var}u_i$ from the compact initial state under heat bath dynamics (x axis rescaled by N^3, y axis by equilibrium value). (b) Rescaling of x axis with an additional logarithm illustrates $O(N^3 \log N)$ time scale. (c) Mixing times for the Markov chains discussed in this Letter. The step ε is uniformly sampled from $(0; 2.5\ell_{\text{free}})$. For the event-chain algorithm, τ_{mix} is measured in lifting moves.

scaling argument for the discrete lowest-k Fourier mode $\bar{U} = (u_1 + u_2 + \cdots + u_{N/2})/N^{3/2}$ yields the $O(N^3)$ behavior in equilibrium: In the limit $N \to \infty$, the standard deviation of $\pi(\bar{U})$ is $O(1)$, and one heat bath step changes $\bar{U}$ by $O(1/N^{3/2})$. A random walk in t yields

$$\frac{1}{N^{3/2}}\sqrt{t} \sim 1 \Rightarrow \tau_{\text{mix}} \gtrsim N^3. \tag{2}$$

The reversible Metropolis algorithm mixes on the same $O(N^3 \log N)$ scale as the reversible heat bath algorithm [see Fig. 1(c)]. At each step, it considers a move from a towards a configuration $\tilde{a}$ with $\tilde{x}_i = x_i + \sigma\varepsilon$, where $\sigma = \pm 1$ samples the forward or backward direction and $\varepsilon > 0$ samples the step from some distribution $p(\varepsilon)$ (we use a uniform distribution on the interval $[0; 2.5\ell_{\text{free}}]$). If $\tilde{a}$ is invalid because of an overlap or an inversion of x_i with x_{i-1} or x_{i+1}, the move $a \to a$ results. The reversible Metropolis algorithm satisfies detailed balance between a and any $\tilde{a}$ simply because the moves $\tilde{a} \to a$ and $a \to \tilde{a}$ are equally likely.

The probability flow $\mathcal{F}_a$ into a has four components for each sphere i (see Fig. 2), namely, accepted forward flow $\mathcal{A}_i^+ = \int d\varepsilon p(\varepsilon)\mathcal{A}_i^+(\varepsilon)$, corresponding to $\sigma = +1$ and analogously accepted backward flow $\mathcal{A}_i^-$ for $\sigma = -1$. Rejected forward and backward flows $\mathcal{R}_i^+$ and $\mathcal{R}_i^-$ from a towards invalid configurations $\tilde{a}$ also contribute to the flow into a. For given ε, these flows

$$\begin{aligned} \mathcal{A}_i^+(\varepsilon) &= \Theta(\delta_i - \varepsilon), & \mathcal{R}_i^+(\varepsilon) &= \Theta(\varepsilon - \delta_{i+1}), \\ \mathcal{A}_i^-(\varepsilon) &= \Theta(\delta_{i+1} - \varepsilon), & \mathcal{R}_i^-(\varepsilon) &= \Theta(\varepsilon - \delta_i), \end{aligned} \tag{3}$$

with Θ the Heaviside step function, are either unity or zero. Moreover, they add up to unity in pairs: $\mathcal{A}_i^+(\varepsilon) + \mathcal{R}_i^-(\varepsilon) = 1$, $\mathcal{A}_i^-(\varepsilon) + \mathcal{R}_i^+(\varepsilon) = 1$, because each move is accepted (or

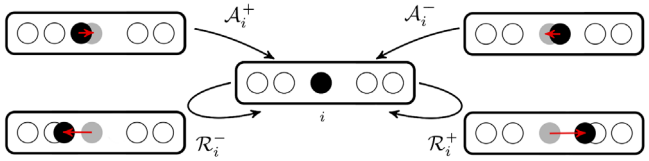

FIG. 2. Metropolis flow into configuration a by moves of sphere i. For given ε, $[\mathcal{A}_i^{\pm}(\varepsilon), \mathcal{R}_i^{\mp}(\varepsilon)] \in \{[1,0],[0,1]\}$ as a move by ε is either accepted or rejected. Flows here are integrated over the step distribution $p(\varepsilon)$. The relation $[\mathcal{A}_i^+(\varepsilon), \mathcal{R}_{i-1}^+(\varepsilon)] \in \{[1,0],[0,1]\}$ justifies the forward Metropolis algorithm.

rejected) under the same condition as its return move. It follows that, for any distribution of ε, the sum of the four flows equals 2.

Global balance requires the total flow $\mathcal{F}_a$ into a valid hard-sphere configuration a to equal $\pi(a) = 1$. For the reversible Metropolis algorithm, there are $2N$ equal choices of the N spheres and two directions $\sigma = \pm 1$, so that

$$\mathcal{F}_a^{\mathrm{rev}} = \frac{1}{2N} \sum_i \underbrace{(\mathcal{A}_i^+ + \mathcal{R}_i^+ + \mathcal{A}_i^- + \mathcal{R}_i^-)}_{=2, \text{ see Eq.}(3)} = 1.$$

Thus, global balance is established, although it was already implied by the detailed-balance condition.

Global balance is also satisfied for the sequential Metropolis algorithm, the historically first irreversible Markov chain [11], which updates spheres sequentially, say, in ascending order in i. At a given time, only a fixed sphere i is updated, and the flow into a configuration a during this move arises from the two choices $\sigma = \pm 1$ for this update of i, which each can be either accepted or rejected

$$\mathcal{F}_a^{\mathrm{seq}} = \frac{1}{2}(\mathcal{A}_i^+ + \mathcal{R}_i^+ + \mathcal{A}_i^- + \mathcal{R}_i^-) = 1. \tag{4}$$

Irreducibility and aperiodicity can also be proven for generic distributions $p(\varepsilon)$. With $p(\varepsilon)$ uniform in the interval $[0, 2.5\ell_{\mathrm{free}}]$, the sequential Metropolis algorithm mixes ~ 1.2 times faster than the reversible Metropolis algorithm, but with the same $O(N^3 \log N)$ scaling.

The relation $\mathcal{A}_i^+(\varepsilon) + \mathcal{R}_i^-(\varepsilon) = 1$ from Eq. (3) can be expressed as $\mathcal{A}_i^+(\varepsilon) + \mathcal{R}_{i-1}^+(\varepsilon) = 1$. This motivates the forward Metropolis algorithm, which attempts, at each time step, a forward move ($\sigma \equiv +1$) sampled from the probability distribution $p(\varepsilon)$, for a randomly sampled sphere i. There are now N equal choices for the moves, and the incoming flow into a configuration a is given by

$$\mathcal{F}_a^{\mathrm{forw}} = \frac{1}{N} \sum_i \underbrace{(\mathcal{A}_i^+ + \mathcal{R}_{i-1}^+)}_{=1} = 1, \tag{5}$$

which, again, establishes global balance. In contrast, the sequential forward Metropolis algorithm violates global balance. In this algorithm, at a given time step, a fixed sphere i is updated, and the flow $\mathcal{F}_a^{\mathrm{seq-forw}}$ into a configuration a, at this time step, arises for a given value ε from a single possible move. The total flow is $\mathcal{F}_a^{\mathrm{seq-forw}} = \mathcal{A}_i^+ + \mathcal{R}_i^+ \neq 1$, so that the sequential forward algorithm is not correct.

Remarkably, we find that the forward Metropolis algorithm mixes on a time scale $O(N^{5/2})$ [see Fig. 1(c)]. The same $N^{5/2}$ time scale also governs the relaxation of the TASEP, a lattice transport model which converges to equilibrium under periodic boundary conditions [18]. An individual TASEP step attempts to move a randomly sampled sphere one site to the right [Fig. 3(a)]. Indeed, the TASEP agrees with the forward Metropolis algorithm restricted to integer x_i and L and steps $\varepsilon \equiv 1$. By tracking the lattice equivalent of $\mathrm{Var} u_i$, we recover the $O(N^{5/2})$ mixing for the TASEP [18], while the symmetric SEP, itself the lattice version of the reversible Metropolis algorithm, mixes in $\sim N^3 \log N$ [26].

The relation $\mathcal{A}_i^+(\varepsilon) + \mathcal{R}_{i-1}^+(\varepsilon) = 1$, for any individual i [see Eq. (6)] provides the motivation for the lifted Metropolis algorithm. Here, moves are attempted in the forward direction $\sigma \equiv +1$, but the active sphere i at time step $t+1$ is determined from the outcome at time step t [see Fig. 3(b) for the discretized example]: As long as sphere i can move, it remains active for the next step. Only when it cannot move, a lifting move $i \to i+1$ takes place, instead, and $i+1$ becomes the active sphere (the physical configuration a does not change during this step). Each configuration is now characterized by the active particle i, in addition to a. The incoming flow $\mathcal{F}_{(a,i)}^{\mathrm{lift}}$ into a lifted configuration (a, i) is either due to an accepted move of i or a rejected move of $i-1$, so that

$$\mathcal{F}_{(a,i)}^{\mathrm{lift}} = \mathcal{A}_i^+ + \mathcal{R}_{i-1}^+ = 1. \tag{6}$$

Global balance again holds. We find that the lifted Metropolis algorithm, run as a Markov chain without restarts (see below) mixes in $O(N^{5/2})$ steps [see Fig. 1(c)]. Thus, it belongs to the TASEP universality class.

Balance conditions relate the stationary probability distribution π at time step t to the distribution at time step

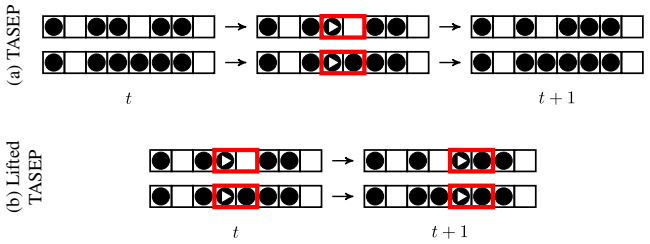

FIG. 3. Update rules for discrete models. (a) The TASEP advances a random sphere, if possible, and it mixes in $O(N^{5/2})$. (b) The lifted TASEP, without restarts, is deterministic. With restarts, it mixes in $O(N^2 \log N)$.

$t+1$, which must agree. For the reversible Metropolis algorithm, as discussed, this condition is satisfied for any sequence of i [see Eq. (5)]. If run for a finite number of steps, the lifted Metropolis algorithm is correct only if started from a random position i. It is advantageous to restart this algorithm after $\lambda \sim N$ time steps by resampling the active sphere i. The chain length λ could also be random, sampled from an appropriate distribution; see the Supplemental Material Section IV [15] for details. We then observe mixing on a time scale $O(N^2 \log N)$, much faster than all previous Markov chains [see Fig. 1(c)]. The $O(N^2)$ time scale is again brought out by a scaling argument for the discrete Fourier mode $\bar{U}$: One chain (sequence of $\sim N$ moves between restarts) of the lifted Metropolis algorithm can change $\bar{U}$ by $O(1/N^{1/2})$. This change is of random sign. A random walk in t/N then yields $\tau_{\mathrm{mix}} \gtrsim N^2$. Indeed, simulations show that restarts every $\lambda \sim N$ time steps yield the fastest mixing [see Fig. 4(a)] and outpace restarts with other scalings with N [see Fig. 4(b), and below].

The lifted Metropolis algorithm also has a discrete counterpart in the lifted TASEP [see Fig. 3(b)]: A single sphere (occupied lattice site) is active and attempts to advance in a forward direction. The sphere remains active if its move is accepted. Otherwise, the lifting index advances to the right-hand neighbor site. The lifted TASEP satisfies global balance, but, without restarts, fails to be irreducible. With restarts every $O(N)$ steps, the lifted TASEP also mixes on an $O(N^2 \log N)$ time scale. The infinitesimal limit of the lifted Metropolis algorithm, $\varepsilon \to 0$, is the ECMC, which also mixes in $O(N^2 \log N)$ lifting moves [see Fig. 1(c)]. Thus, it also belongs to the lifted TASEP universality class.

The dynamical universality classes for local 1D hard-sphere algorithms are summarized in Table I. In the lifted TASEP universality class, each sphere only moves $O(N \log N)$ times to reach equilibrium, almost saturating the lower bound, as $O(N)$ are required to detach each sphere from the compact initial state.

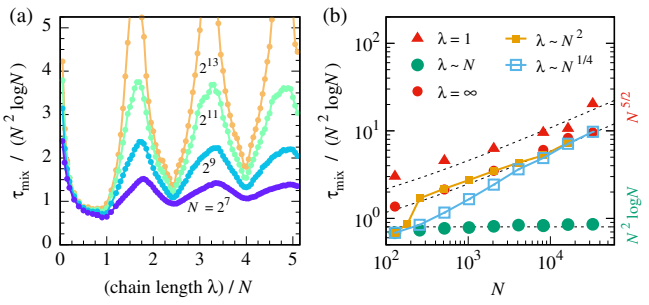

FIG. 4. Restarts in the lifted Metropolis algorithm. (a) Mixing time as a function of chain length $\lambda \sim N$: The $O(N^2 \log N)$ time scale is preserved but mixing is fastest for $\lambda \simeq 0.9N$. (b) $O(N^{5/2})$ mixing for $\lambda \sim N^{1/4}$ [asymptotically equivalent to the (unlifted) forward Metropolis algorithm, $\lambda = 1$, red triangles] and for $\lambda \sim N^2$ (asymptotically equivalent to lifted Metropolis without any restarts, $\lambda = \infty$, red bullets).

The irreversible Markov chains presented here are best viewed from their deterministic roots, both on the lattice and in the continuum. Indeed, the lifted TASEP without restarts is a deterministic lattice-gas automaton satisfying global balance, rather than a Markov chain [see Fig. 3(b)]. Irreducibility and aperiodicity call for an element of randomness that can be supplied by restarts. Here, $\lambda \sim N$ (lifted TASEP) and $\lambda \equiv 1$ (TASEP) represent different universality classes. In the continuum, the deterministic root is an algorithm without restarts and invariant step ε, which satisfies global balance. All presented Metropolis-type Markov chains may be obtained from this root by resampling of i or ε. Algorithms belong to different universality classes, depending on the resampling rate: Resampling ε at every step leads to the lifted Metropolis algorithm without restarts which is in the TASEP class. The additional resampling of i every $\lambda \sim N$ time steps yields the lifted Metropolis algorithm with restarts, which is in the lifted TASEP class [see Fig. 4(a), the oscillations suggest that resamplings should ideally correspond to $\sim N/2$ successfully moved spheres and that multiples of N should be avoided]. More infrequent restarts ($\lambda \sim N^2$) or more frequent ones ($\lambda \sim N^{1/4}$) lead back to the nonlifted TASEP class [see Fig. 4(b)]. Numerical computations of the asymptotic mixing time scales for $\lambda \sim N^\alpha$ with α only slightly different from 1 will require very large system sizes in order to overcome the oscillations at $\alpha = 1$.

Injecting randomness with a rate $\sim 1/N$ is, thus, optimal for mixing. A particular limit of the continuum root algorithm is the ECMC without restarts. In 1D, it agrees with Newtonian dynamics with a single sphere of nonzero velocity, and never mixes. Resampling i with rate $\sim 1/N$ (corresponding to an occasional restart of Newtonian dynamics) turns the deterministic dynamics into a very fast local algorithm, the ECMC, which is in the lifted TASEP class, and mixes in $O(N^2 \log N)$ lifting moves [see Fig. 1(c)].

The infinitesimal displacements of the ECMC are crucial for its generalization to situations where particles simultaneously interact with several others, such as in more than 1D, or due to long-range forces. This generalization calls

TABLE I. Mixing time scales for local 1D hard-sphere algorithms on the continuum and on the lattice. The Markov chains on the lowest row all incorporate restarts. See the Supplemental Material Section IV [15] for pseudocode implementations.

Local 1D hard-sphere Markov chains		Mixing time scale
Continuous	Discrete	
Heat bath [14], Metropolis	Symmetric SEP [26]	$N^3 \log N$
Forward and lifted Metropolis without restarts	TASEP [18]	$N^{5/2}$
Event-chain, lifted Metropolis	Lifted TASEP	$N^2 \log N$

for the factorized Metropolis algorithm [22,27,28]. The finite-step lifted Metropolis algorithm remains correct for repulsive interactions restricted to nearest neighbors. Indeed, soft repulsive particles with a $1/x^{12}$ potential reproduce the mixing behavior of hard spheres, with $O(N^2 \log N)$ mixing for the lifted Metropolis and $O(N^{5/2})$ mixing for lifted Metropolis without restarts (see the Supplemental Material Section V [15] for details). The ECMC has been applied successfully (see Ref. [23] for a review), but prior to the present Letter, its mixing behavior was not characterized in detail, beyond some partial evidence for faster mixing time scales [24].

In the future, it will be important to clarify how the different universality classes identified in the present Letter carry over to higher dimensions and to what degree they depend on the asymmetry of the step distribution. Exact solutions of some of the models in Table I may be possible. This would help establish the conceptual framework of lifting and of irreversible Markov chains beyond the single-particle level [21]. More generally, the mixing dynamics of hard spheres from an initial compact state may be interpreted as the equilibration process of a physical system in response to a sudden change in its Hamiltonian at time $t = 0$. The analysis of the asymmetric evolution of shock fronts may shed further light on the new universality class, as previously for the well-studied TASEP class. An example for a system strongly out of equilibrium, it will be interesting to study how entropy production differs between the three universality classes, and how they may be integrated within the nonequilibrium fluctuation theorems [29,30].

S. C. K. thanks the Institut Philippe Meyer for support, and T. Franosch for helpful discussions. We thank the referees for advice on an earlier version of the article. Publication was supported by the Emerging Talents Intitiative of the FAU Erlangen-Nuremberg.

*sebastian.kapfer@fau.de
†werner.krauth@ens.fr

[1] J. G. Kirkwood and E. Monroe, J. Chem. Phys. **9**, 514 (1941); W. G. Hoover and F. H. Ree, J. Chem. Phys. **49**, 3609 (1968).
[2] B. J. Alder and T. E. Wainwright, Phys. Rev. **127**, 359 (1962).
[3] E. P. Bernard and W. Krauth, Phys. Rev. Lett. **107**, 155704 (2011).
[4] S. Asakura and F. Oosawa, J. Chem. Phys. **22**, 1255 (1954).
[5] W. Krauth, *Statistical Mechanics: Algorithms and Computations* (Oxford University Press, New York, 2006).
[6] B. I. Halperin and D. R. Nelson, Phys. Rev. Lett. **41**, 121 (1978); A. P. Young, Phys. Rev. B **19**, 1855 (1979).
[7] B. J. Alder and T. E. Wainwright, J. Chem. Phys. **27**, 1208 (1957).
[8] B. J. Alder and T. E. Wainwright, Phys. Rev. A **1**, 18 (1970).
[9] G. Parisi and F. Zamponi, Rev. Mod. Phys. **82**, 789 (2010).
[10] S. Torquato and F. H. Stillinger, Rev. Mod. Phys. **82**, 2633 (2010).
[11] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).
[12] R. Kannan, M. W. Mahoney, and R. Montenegro, in *Algorithms and Computation: 14th International Symposium, ISAAC 2003, Kyoto, Japan, 2003: Proceedings,* edited by T. Ibaraki, N. Katoh, and H. Ono, Lecture Notes in Computer Science, Vol. 2906 (Springer, Berlin, 2003), pp. 663–675.
[13] L. Tonks, Phys. Rev. **50**, 955 (1936).
[14] D. Randall and P. Winkler, in *Approximation, Randomization and Combinatorial Optimization*, edited by C. Chekuri *et al.*, Lecture Notes in Computer Science, Vol. 3624 (Springer, Berlin, 2005).
[15] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.119.240603 for further details on the convergence criteria and mixing times of general Markov chains, and for more information on the discussed algorithms, including their density dependence, full pseudocode implementations, and an extension of the finite-step lifted Metropolis to soft potentials.
[16] L.-H. Gwa and H. Spohn, Phys. Rev. Lett. **68**, 725 (1992).
[17] T. Chou, K. Mallick, and R. K. P. Zia, Rep. Prog. Phys. **74**, 116601 (2011).
[18] J. Baik and Z. Liu, J. Stat. Phys. **165**, 1051 (2016).
[19] E. P. Bernard, W. Krauth, and D. B. Wilson, Phys. Rev. E **80**, 056704 (2009).
[20] F. Chen, L. Lovász, and I. Pak, *Proceedings of the 31st Annual ACM Symposium on Theory of Computing* (ACM, New York, 1999), p. 275.
[21] P. Diaconis, S. Holmes, and R. M. Neal, Ann. Appl. Probab. **10**, 726 (2000).
[22] M. Michel, S. C. Kapfer, and W. Krauth, J. Chem. Phys. **140**, 054116 (2014).
[23] S. C. Kapfer and W. Krauth (to be published).
[24] Y. Nishikawa, M. Michel, W. Krauth, and K. Hukushima, Phys. Rev. E **92**, 063306 (2015).
[25] D. Aldous and P. Diaconis, Am. Math. Mon. **93**, 333 (1986).
[26] H. Lacoin, Ann. Inst. H. Poincaré Probab. Statist. **53**, 1402 (2017).
[27] E. A. J. F. Peters and G. de With, Phys. Rev. E **85**, 026703 (2012).
[28] S. C. Kapfer and W. Krauth, Phys. Rev. E **94**, 031302 (2016).
[29] D. J. Evans and D. J. Searles, Adv. Phys. **51**, 1529 (2002).
[30] U. Seifert, Rep. Prog. Phys. **75**, 126001 (2012).